

ORACLE CONFIDENTIAL

Patent Application for:
CODE GENERATION FACILITY

Inventors:
Jonathan Maron

Oracle Corporation

CODE GENERATION FACILITY

Field of the Invention

5 The present invention relates to a method, system, and computer program product for generating code from compiled code deployed to an application server.

Background of the Invention

10 Many Java technologies currently require the generation of classes based on some input class. A prime example is Enterprise Java Beans (EJBs), which require the generation of a particular application server's container classes to wrap the developer supplied bean class.

 To date this process of code generation has largely been undertaken with
15 invocations of methods that use printout statements to build the required files, generally one file at a time. In addition, most of these generation facilities attempt to create a complete file at once rather than taking a more modular approach to the assembly of the generated file. Although this process works for small projects, it quickly becomes unwieldy for larger projects, leading to
20 increased expense. In addition, tracking down errors in code generation or

making modifications to accommodate new requirements is a difficult, expensive, and time-consuming process since the correct method and the correct print statement needs to be found amongst many similar looking statements. A need arises for a technique by which code can be generated that reduces the cost
5 for large projects and that provides easier and quicker ways of finding errors and making modifications.

Summary of the Invention

The present invention is a system and method for generating code that
10 reduces the cost for large projects and that provides easier and quicker ways of finding errors and making modifications. Whereas the prior art techniques were rather unwieldy and cumbersome, the present invention provides a modular approach to code generation - event handlers handle precise segments of the code generation process rather than attempting to create a complete file. Thus, the
15 whole code generation engine within an application becomes more maintainable and reliable, reducing cost and time spent.

A more maintainable and reliable code generation facility provides the capability to significantly decrease any generation errors that may occur and that may ultimately lead to customer dissatisfaction. In addition, the ability to quickly

and effectively modify the code generation process will allow the server to adapt to new generation requirements for particular technologies quickly, thus significantly cutting down on time-to-market for new generation dependent features.

5 In addition the present invention has potential applicability in many code generation arenas such as CORBA ORBs, XML code bindings such as JAXB, etc.

In one embodiment of the present invention, a method of generating code to be deployed in an application server comprises the steps of receiving a
10 compiled file to be deployed, introspecting this input class to generate information relating to the input class, generating a markup-language description of the input class based on the generated information relating to the input class, creating an event handler for a method node found in the markup-language description, registering the event handler, parsing the markup-language
15 description and invoking the registered event handler, and generating output code using the invoked event handler.

In one aspect of the present invention, the archive file is a compiled Enterprise Java Bean file. The step of introspecting an input class included in the archive file may comprise the steps of extracting information identifying methods

included in the input class and for each method, extracting information relating to parameters of the method and exceptions generated by the method. The step of generating a markup-language description of the input class may comprise the step of generating an Extensible Markup-Language description of the input class

5 based on the generated information relating to the input class. The step of creating an event handler may comprise the step of creating a programming language implementation of a Simple Application Programming Interface for Extensible Markup Language (SAX) event handler for a method node found in the class' XML description. The step of parsing the XML description and

10 invoking the registered event handlers may comprise the step of parsing the markup language description using a SAX parser and invoking the registered SAX event handlers.

In one aspect of the present invention, the step of creating an event handler comprises the step of creating a plurality of event handlers for a method

15 node found in the markup language description. The step of registering the event handler may comprise the step of registering each of the plurality of event handlers. The step of parsing the markup language description and invoking the registered event handler may comprise the step of parsing the markup language description and invoking each of the plurality of registered event handlers. The

step of generating output code may comprise the step of generating output code using each of the plurality of invoked event handlers in parallel.

Brief Description of the Drawings

5 The details of the present invention, both as to its structure and operation, can best be understood by referring to the accompanying drawings, in which like reference numbers and designations refer to like elements.

Fig. 1 is an exemplary block diagram of an application server, in which the present invention may be implemented.

10 Fig. 2 is an exemplary flow diagram of a process for generating code according to the present invention.

Fig. 3 is an exemplary block diagram of a computer system, in which an application server, such as that shown in Fig. 1, may be implemented.

15 Fig. 4 is an exemplary architecture of a SAX implementation that may be used by the present invention.

Detailed Description of the Invention

The present invention is a system and method for generating code that reduces the cost for large projects and that provides easier and quicker ways of

ORACLE CONFIDENTIAL

finding errors and making modifications. In a preferred embodiment, the present invention relies on the Simple API for XML (SAX) to drive the generation of code. SAX is an event-driven, serial-access mechanism for accessing XML documents. SAX was the first widely adopted API for XML in Java, and is a “de facto” standard. A typical version of SAX contains a number of core classes and interfaces together with optional helper classes and demonstration classes. SAX is described in greater detail below.

An exemplary block diagram of an application server 100, in which the present invention may be implemented, is shown in Fig. 1. The exemplary application server shown in Fig. 1 is based on the well-known ORACLE® application server. However, this application server is merely an example. The present invention is applicable to other application servers as well, such as those from IBM®, MICROSOFT®, etc.

An Application Server generally interacts with both clients 102 and back-end EIS systems 104. The application server provides an Enterprise Java Bean (EJB) container 106 to facilitate access to EIS and/or database systems 104. A container is a controlled run-time environment for an application component (e.g., an EJB 108) that also provides an interface to components

ORACLE CONFIDENTIAL

executing within the container. Also, containers provide basic management functionality such as life cycle management of the components, security, deployment, threading, etc. An EJB 108 is a software component that implements the interfaces defined by its respective container. The interfaces of
5 the EJB 108 are defined in a manner that permits the functionality of the respective EJB 108 to be accessed in a modular and dynamic manner.

The application server 100 also provides a Web container 110. The Web container 110 may provide a controlled run-time environment for components that interact with a client. For example, servlet components 112
10 (persistent Web server processes that are capable of processing multiple requests) may be implemented to dynamically create hypertext markup language (HTML) or XML responses to requests from clients. Other suitable components may be implemented within the Web container.

In general, the middle tier may be implemented utilizing “middleware servers” or “application servers.” An application server typically includes many dissimilar services that are appropriate to create a stable and robust environment for enterprise applications. A service represents a predefined task that is provided by a piece of software in a well defined and predictable manner.
15 Typically, an application server may comprise services that implement

management functionality, such as MBean server 114, in which management services may be implemented as MBeans 116. The management services of the application server may be responsible for starting, registering, monitoring, and stopping services. Management services may perform other tasks such as

5 thread pooling, security management, state management, class loading, load-balancing, dynamic application launching, and/or the like. Secondly, an application server may provide a basic set of services that may be commonly utilized by a wide range of enterprise applications (e.g., hypertext transfer protocol (HTTP) processing). Third, an application server may comprise

10 application specific services that implement the business or other logic of a given enterprise application.

An exemplary flow diagram of a process 200 for generating code according to the present invention is shown in Fig. 2. Typically, a bean archive file is deployed to an application server. For example, if the bean archive file

15 contains an Enterprise Java Bean implementation, the archive file includes one or more classes of objects. Process 200 begins with step 202, in which the input class, that is the object class included in the archive file, is introspected. Introspection involves extracting information relating to the objects, methods, etc., that are included in the class. For example, information identifying

methods included in the input class may be extracted and for each method, information relating to parameters of the method are extracted. In Java, introspection may be performed using the Java Reflection API.

In step 204, an XML description of the Java bean class is generated based

- 5 on the class information gathered via the introspection performed in step 202.

For example,

```
<bean name="MyBean" package="MyPackage">  
  <method name="method1">  
    <param type="java.lang.String" name="param1"/>  
    <exception type="MyPackage.MyException"/>  
    <return type="void"/>  
  </method>  
</bean>
```

15

In step 206, SAX event handlers are created for the appropriate nodes in the XML code. For example, in the example above, a <method> node handler would be at least one of the handlers registered. These event handlers can then drive the generation of the code by analyzing the information of the node and any

relevant child nodes. Based on this analysis the code can drive the generation process in any way a developer would see fit, such as:

- Invocations of print statements
- Templating technologies such as XSL.

5 In step 208, the event handlers are registered with a SAX parser and, in step 210, the parsing of the Java bean class XML descriptor commences. As the event handlers are invoked by the parser, step 212, the code generation proceeds. Note that this facility could potentially be leveraged to generate more than one file in parallel if, for example, an event handler leverages multiple templates. For
10 example, if a method affects multiple files, all handlers for a method may be invoked at once and the code for each file generated in parallel.

An exemplary block diagram of a computer system 300, in which an application server, such as that shown in Fig. 1, may be implemented, is shown in Fig. 3. System 300 is typically a programmed general-purpose computer system, such as a personal computer, workstation, server system, and minicomputer or mainframe computer. System 300 includes one or more processors (CPUs) 302A-302N, input/output circuitry 304, network adapter 306, and memory 308. CPUs 302A-302N execute program instructions in order to carry out the functions of the present invention. Typically, CPUs
15

ORACLE CONFIDENTIAL

302A-302N are one or more microprocessors, such as an INTEL PENTIUM® processor. Fig. 3 illustrates an embodiment in which System 300 is implemented as a single multi-processor computer system, in which multiple processors 302A-302N share system resources, such as memory 308, 5 input/output circuitry 304, and network adapter 306. However, the present invention also contemplates embodiments in which System 300 is implemented as a plurality of networked computer systems, which may be single-processor computer systems, multi-processor computer systems, or a mix thereof.

Input/output circuitry 304 provides the capability to input data to, or 10 output data from, database/System 300. For example, input/output circuitry may include input devices, such as keyboards, mice, touchpads, trackballs, scanners, etc., output devices, such as video adapters, monitors, printers, etc., and input/output devices, such as, modems, etc. Network adapter 306 interfaces database/System 300 with Internet/intranet 310. Internet/intranet 310 15 may include one or more standard local area network (LAN) or wide area network (WAN), such as Ethernet, Token Ring, the Internet, or a private or proprietary LAN/WAN.

Memory 308 stores program instructions that are executed by, and data that are used and processed by, CPU 302 to perform the functions of system

300. Memory 308 may include electronic memory devices, such as random-access memory (RAM), read-only memory (ROM), programmable read-only memory (PROM), electrically erasable programmable read-only memory (EEPROM), flash memory, etc., and electro-mechanical memory, such as

5 magnetic disk drives, tape drives, optical disk drives, etc., which may use an integrated drive electronics (IDE) interface, or a variation or enhancement thereof, such as enhanced IDE (EIDE) or ultra direct memory access (UDMA), or a small computer system interface (SCSI) based interface, or a variation or enhancement thereof, such as fast-SCSI, wide-SCSI, fast and wide-SCSI, etc,

10 or a fiber channel-arbitrated loop (FC-AL) interface.

In the example shown in Fig. 3, memory 308 includes HTTP listener layer 108, virtual path manager 109, application server layer 110, applications layer 112, ORB 114, code generation objects 310, SAX objects 312, and operating system 314. HTTP listener layer 108 is made up of listeners, the
15 adapter interface, and dispatchers. Listeners are HTTP servers; they handle incoming requests and route them to the dispatcher. The dispatcher forwards requests to the virtual path manager 109. The virtual path manager maps a request to a cartridge type and passes this information back to the dispatcher. The virtual path manager also passes back authentication requirements to the

dispatcher. The Application Server layer 110 provides resource management in handling requests for applications deployed as cartridges on the server. It provides a common set of components for managing these applications. These components include load balancing, logging, automatic failure recovery, security, directory, and transaction components. The Applications layer 112 is made up of applications, cartridges, and cartridge servers. Applications and cartridges are the two main objects that you use when building applications for the application server environment. ORB 114 acts as the middleware between clients and servers. Code generation objects 310 perform the steps of and store the data for code generation process 200, shown in Fig. 2. SAX objects 312 implement the SAX functionality, described below. Operating system 314 provides overall system functionality.

As shown in Fig. 3, the present invention contemplates implementation on a system or systems that provide multi-processor, multi-tasking, multi-process, and/or multi-thread computing, as well as implementation on systems that provide only single processor, single thread computing. Multi-processor computing involves performing computing using more than one processor. Multi-tasking computing involves performing computing using more than one operating system task. A task is an operating system concept that refers to the

combination of a program being executed and bookkeeping information used by the operating system. Whenever a program is executed, the operating system creates a new task for it. The task is like an envelope for the program in that it identifies the program with a task number and attaches other
5 bookkeeping information to it. Many operating systems, including UNIX®, OS/2®, and WINDOWS®, are capable of running many tasks at the same time and are called multitasking operating systems. Multi-tasking is the ability of an operating system to execute more than one executable at the same time. Each executable is running in its own address space, meaning that the executables
10 have no way to share any of their memory. This has advantages, because it is impossible for any program to damage the execution of any of the other programs running on the system. However, the programs have no way to exchange any information except through the operating system (or by reading files stored on the file system). Multi-process computing is similar to multi-
15 tasking computing, as the terms task and process are often used interchangeably, although some operating systems make a distinction between the two.

SAX is an event-driven, serial-access mechanism for accessing XML documents. SAX was the first widely adopted API for XML in Java, and is a “de

ORACLE CONFIDENTIAL

facto” standard. A typical version of SAX contains a number of core classes and interfaces together with optional helper classes and demonstration classes.

An exemplary architecture of a SAX implementation 400 is shown in Fig.

4. The parser 402 wraps a SAXReader object 404. When the parser's 402
- 5 parse() method is invoked, the reader 404 invokes one of several callback methods implemented in the application. Those methods are defined by the interfaces ContentHandler 406, ErrorHandler 408, DTDHandler 410, and EntityResolver 412.

10 Here is a summary of the key SAX APIs:

SAXParserFactory 414

A SAXParserFactory 414 object creates an instance of the parser determined by the system property,
15 javax.xml.parsers.SAXParserFactory.

SAXParser 402

The SAXParser interface 402 defines several kinds of parse()
methods. In general, you pass an XML data source 416 and a

`DefaultHandler` object to the parser, which processes the XML 416 and invokes the appropriate methods in the handler object.

SAXReader 404

5 The `SAXParser` 404 wraps a `SAXReader` 402. Typically, you
don't care about that, but every once in a while you need to get hold of it using
`SAXParser's getXMLReader()`, so you can configure it. It is the
`SAXReader` 402 that carries on the conversation with the SAX event
handlers you define.

10

DefaultHandler

Not shown in Fig. 4, a `DefaultHandler` implements the
`ContentHandler` 406, `ErrorHandler` 408, `DTDHandler` 410, and
`EntityResolver` 412 interfaces (with null methods), so you can override
15 only the ones you're interested in.

ContentHandler 406

Methods like `startDocument`, `endDocument`, `startElement`,
and `endElement` are invoked when an XML tag is recognized. This

interface also defines methods characters and processingInstruction, which are invoked when the parser encounters the text in an XML element or an inline processing instruction, respectively.

5 **ErrorHandler** 408

Methods `error`, `fatalError`, and `warning` are invoked in response to various parsing errors. The default error handler throws an exception for fatal errors and ignores other errors (including validation errors).

That's one reason you need to know something about the SAX parser, even if

- 10 you are using the DOM. Sometimes, the application may be able to recover from a validation error. Other times, it may need to generate an exception. To ensure the correct handling, you'll need to supply your own error handler to the parser.

15 **DTDHandler** 410

Defines methods you will generally never be called upon to use. Used when processing a DTD to recognize and act on declarations for an *unparsed entity*.

EntityResolver 412

The `resolveEntity` method is invoked when the parser must identify data identified by a URI. In most cases, a URI is simply a URL, which specifies the location of a document, but in some cases the document may be

5 identified by a URN--a *public identifier*, or name, that is unique in the Web space. The public identifier may be specified in addition to the URL. The `EntityResolver` can then use the public identifier instead of the URL to find the document, for example to access a local copy of the document if one exists. A typical application implements most of the `ContentHandler`

10 methods, at a minimum. Since the default implementations of the interfaces ignore all inputs except for fatal errors, a robust implementation may want to implement the `ErrorHandler` methods, as well.

These SAX classes and interfaces fall into groups:

15 • Interfaces implemented by the parser:

- Parser and AttributeList (required), and Locator (optional)
- Interfaces implemented by the application:

- DocumentHandler, ErrorHandler, DTDHandler, and EntityResolver
(all optional: DocumentHandler will be the most important one for typical XML applications)
- Standard SAX classes:
 - InputSource, SAXException, SAXParseException, HandlerBase
- Optional Java-specific helper classes:
 - ParserFactory, AttributeListImpl, and LocatorImpl
- Java demonstration classes in the nul package:
 - SystemIdDemo, ByteStreamDemo, CharacterStreamDemo, and EntityDemo

10 Interfaces for Parser Writers:

A SAX-conformant XML parser needs to implement only two or three simple interfaces; in fact, it is even possible (and quite common) to implement all 15 of these interfaces in a single class, if desired.

Parser

This is the main interface to a SAX parser: it allows the user to register handlers for callbacks, to set the locale for error reporting, and to start an XML parse.

AttributeList

This simple interface allows users to iterate through an attribute list -- the parser can implement it in the same class as the SAX driver, or implement it in a separate class (it does not need to be persistent). There is a convenience 5 implementation available in the org.xml.sax.helpers.AttributeListImpl.

Locator

This simple interface allows users to find the current location in the XML source document -- the parser can implement it in the same class as the driver, or implement it in a second class (it does not need to be persistent).

10

Interfaces for Application Writers

A SAX application may implement any or none of the following interfaces, as required (simple XML applications may need only DocumentHandler and possibly ErrorHandler). An application can implement all 15 of these interfaces in a single class, if desired.

DocumentHandler

This is the interface that applications will probably use the most -- in many cases, it is the only one that they will need to implement. If an application provides an implementation of this interface, it will receive notification of basic

ORACLE CONFIDENTIAL

document-related events like the start and end of elements. It is this interface that will be implemented to drive the code generation in this invention.

ErrorHandler

If an application needs to use special error handling, then it must provide

- 5 an implementation of this interface.

DTDHandler

If an application needs to work with notations and unparsed (binary) entities, it must implement this interface to receive notification of the NOTATION and ENTITY declarations.

10 EntityResolver

If an application needs to do redirection of URIs in documents (or other types of custom handling), it must provide an implementation of this interface.

Standard SAX Classes

15 SAX consists mostly of interfaces rather than classes, but the interfaces refer to two standard exception classes, and a third is provided for universal convenience. These classes are useful for both parser and application writers.

InputSource

This class contains all of the necessary information for a single input source, including a public identifier, system identifier, byte stream, and character stream (as appropriate). The application must instantiate at least one InputSource for the Parser, and the EntityHandler may instantiate others.

5 SAXException

This class represents a general SAX exception.

SAXParseException

This class represents a SAX exception tied to a specific point in an XML source document.

10 HandlerBase

This class provides default implementations for DocumentHandler, ErrorHandler, DTDHandler, and EntityResolver: parser writers can use this to provide a default implementation when the user does not specify handlers, and application writers can subclass this to simplify handler writing.

15

Java-Specific Helper Classes

These classes are not part of the core SAX distribution, and may not be available in SAX implementations in other languages: they are provided simply as a convenience for Java programmers.

ParserFactory

An application can use the static methods in this class to load SAX parsers dynamically at run time, based on the class name.

AttributeListImpl

5 An application can use this convenience class to make a persistent copy of an AttributeList, or parser can use it to supply a default implementation of AttributeList to the application.

LocatorImpl

An application can use this convenience class to make a persistent 10 snapshot of a Locator's values at a specific point in the parse.

It is important to note that while the present invention has been described in the context of a fully functioning data processing system, those of ordinary skill in the art will appreciate that the processes of the present invention are 15 capable of being distributed in the form of a computer readable medium of instructions and a variety of forms and that the present invention applies equally regardless of the particular type of signal bearing media actually used to carry out the distribution. Examples of computer readable media include recordable-type

ORACLE CONFIDENTIAL

media such as floppy disc, a hard disk drive, RAM, and CD-ROM's, as well as transmission-type media, such as digital and analog communications links.

Although specific embodiments of the present invention have been described, it will be understood by those of skill in the art that there are other 5 embodiments that are equivalent to the described embodiments. Accordingly, it is to be understood that the invention is not to be limited by the specific illustrated embodiments, but only by the scope of the appended claims.